

# Annotation Inconsistency and Entity Bias in MultiWOZ

Kun Qian<sup>†\*</sup>, Ahmad Berrami<sup>‡</sup>, Zhouhan Lin<sup>‡</sup>, Ankita De<sup>‡</sup>, Alborz Geramifard<sup>‡</sup>,  
Zhou Yu<sup>†\*</sup>, Chinnadhurai Sankar<sup>‡</sup>

<sup>†</sup>Columbia University

{kq2157, zy2461}@columbia.edu

<sup>‡</sup>Facebook AI

{beirami, deankita, alborzgeramifard, chinnadhurai}@fb.com

## Abstract

MultiWOZ (Budzianowski et al., 2018) is one of the most popular multi-domain task-oriented dialog datasets, containing 10K+ annotated dialogs covering eight domains. It has been widely accepted as a benchmark for various dialog tasks, e.g., dialog state tracking (DST), natural language generation (NLG) and end-to-end (E2E) dialog modeling. In this work, we identify an overlooked issue with dialog state annotation inconsistencies in the dataset, where a slot type is tagged inconsistently across similar dialogs leading to confusion for DST modeling. We propose an automated correction for this issue, which is present in 70% of the dialogs. Additionally, we notice that there is significant entity bias in the dataset (e.g., “cambridge” appears in 50% of the destination cities in the train domain). The entity bias can potentially lead to named entity memorization in generative models, which may go unnoticed as the test set suffers from a similar entity bias as well. We release a new test set with all entities replaced with unseen entities. Finally, we benchmark joint goal accuracy (JGA) of the state-of-the-art DST baselines on these modified versions of the data. Our experiments show that the annotation inconsistency corrections lead to 7-10% improvement in JGA. On the other hand, we observe a 29% drop in JGA when models are evaluated on the new test set with unseen entities.

## 1 Introduction

Commercial virtual assistants are used by millions via devices such as Amazon Alexa, Google Assistant, Apple Siri, and Facebook Portal. Modeling such conversations requires access to high quality and large task-oriented dialog datasets. Many

researchers have devoted great efforts to creating such datasets and multiple task-oriented dialog datasets, e.g., WOZ (Rojas-Barahona et al., 2017), MultiWOZ (Budzianowski et al., 2018), TaskMaster (Byrne et al., 2019), Schema-Guided Dialog (Rastogi et al., 2019) with fine-grained dialog state annotation have been released in the recent years.

Among task-oriented dialog datasets, MultiWOZ (Budzianowski et al., 2018) has gained the most popularity. The dataset contains 10k+ dialogs and covers eight domains: *Attraction, Bus, Hospital, Hotel, Restaurant, Taxi, Train* and *Police*. Each dialog can cover one or multiple domains. The inclusion of detailed annotations, e.g., task goal, dialog state, and dialog acts for both user side and system side, renders MultiWOZ a universal benchmark for many dialog tasks, such as dialog state tracking (Zhang et al., 2019, 2020a; Heck et al., 2020), dialog policy optimization (yang Wu et al., 2019; Wang et al., 2020a,b) and end-to-end dialog modeling (Zhang et al., 2020b; Hosseini-Asl et al., 2020; Peng et al., 2020). Several recent papers, such as SimpleTOD (Hosseini-Asl et al., 2020), TRADE (Wu et al., 2019), MarCo (Wang et al., 2020b), evaluate their models solely on the MultiWOZ dataset, which makes their findings highly dependent on the quality of this dataset.

Over the last couple of years, several sources of errors have been identified and corrected on MultiWOZ. Wu et al. (2019) pre-processed the dataset by normalizing the text and annotations. Eric et al. (2019) further corrected the dialog state annotations on over 40% dialog turns and proposed MultiWOZ 2.1. Recently, Zang et al. (2020) identified some more error types, fixed annotations from nearly 30% dialogs and added span annotations for both user and system utterances, leading to the MultiWOZ 2.2. Concurrent with this work, Han et al. (2020) released MultiWOZ 2.3 further looking at

\*The work of KQ and ZY was done as a research intern and a visiting research scientist at Facebook AI.

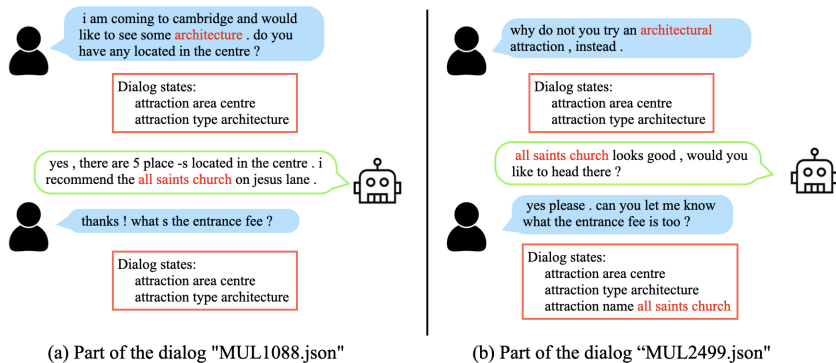


Figure 1: An example of two dialogs with inconsistent annotations. In the left dialog (a), the “*attraction name*” mentioned in the system utterance is not annotated in the dialog state, while in the right dialog (b), the dialog state annotation includes the “*attraction name*”.

annotation consistency by exploring co-references.

While most of the previous works focus on correcting the annotation errors and inconsistencies within a dialog, where annotation contradicts the dialog context, we noticed another overlooked source of confusion for dialog state modeling, namely annotation inconsistency across different dialogs. We first show that the dialogs have been annotated inconsistently with respect to the slot type ‘Name’. Figure 1 shows two dialogs in the *Attraction* domain with similar context. In one dialog the attraction name is annotated while not in the other one. This inconsistency leads to a fundamental confusion for dialog state modeling whether to predict the attraction name or not in similar scenarios. In Section 3, we dive deeper into this problem and propose an automated correction for this problem.

We further found a second source of potential issue, entity bias, where the distribution of the slot value in the dataset is highly imbalanced. In Figure 2, we observed that “cambridge” appears as the train destination in 50% of the dialogs in train domain while there are 13 destinations. As a result, a dialog system trained on this imbalanced data might be more likely to generate “cambridge” as the slot value even though “cambridge” might not even be mentioned in dialog history. In Section 4, we discuss this problem in more detail and suggest a new test set with all entities replaced with ones never seen during training. Finally, in Section 5, we benchmark the state-of-the-art dialog state tracking models on these new versions of data and conclude with our findings.

Our contributions in this paper can be summarized as follows:

- We identify annotation inconsistency across

slot value	Count Num.
<b>cambridge</b>	<b>8,086</b>
london liverpool street	760
leicester	746
stansted airport	711
stevenage	710
ely	695
norwich	692
bishops stortford	667
broxbourne	634
peterborough	630
birmingham new street	624
london kings cross	609
kings lynn	574
total	16,138

Figure 2: The distribution of slot values for slot type “destination” in the “train” domain.

similar dialogs as a new source of error that leads to confusion for DST modeling. We also propose an automated correction for these inconsistencies which result in changes in 66% of the dialogs in the dataset, and release the new training/validation/test data.<sup>1</sup>


- We identify that several slot types suffer from severe entity bias that potentially lead to models memorizing these entities, and release a new test set where all entities are replaced with ones not seen in training data.
- We benchmark state-of-the-art DST models on the new version of data, and observe a 7-10% performance improvement in joint goal accuracy compared to MultiWOZ 2.2. For the data bias, we observe that models evaluated on the new test set with unseen entities suffer from a 29% performance drop potentially caused by memorization of these entities.

## 2 Related Work

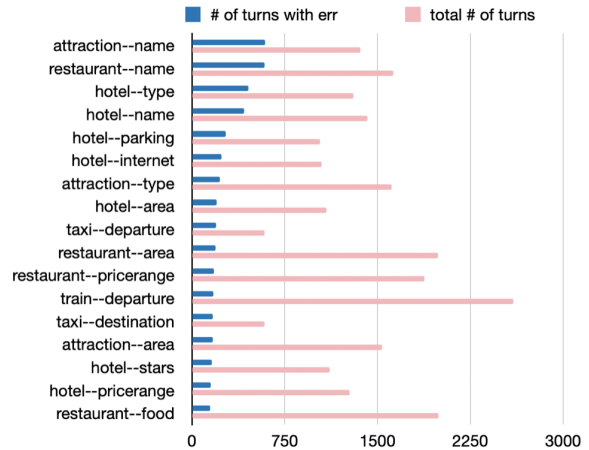
### 2.1 MultiWOZ 2.1

MultiWOZ 2.1 (Eric et al., 2019) mainly focuses on the semantic annotation errors. It identifies five main error types for dialog state annotation: delayed markups, multi-annotations, mis-annotations, typos and forgotten values. The delayed markups refer to the slot values that are annotated one or more turns after where the values show up in the utterances. Multi-annotations mean that multiple slots share the same slot value in a single turn. Mis-annotations represent the errors where a slot value

<sup>1</sup>The code and data is available at [https://github.com/qbetterk/multiwoz/tree/master/data/MultiWOZ\\_2.2%2B](https://github.com/qbetterk/multiwoz/tree/master/data/MultiWOZ_2.2%2B)



(a) The generated result for dialog "MUL1088.json"



(b) Distribution of turn numbers with err over slot types

Figure 3: (a) The generated dialog state results (in green rectangles) using SimpleTOD. The model generates the “*attraction name all saints church*” in the second turn; (b) The distribution of the number of turns where SimpleTOD makes a mistake. The slot type “*name*” from domains “*attraction*”, “*hotel*” and “*restaurant*” and slot type “*type*” from “*hotel*” domain have more error turns than others.

is assigned to a wrong slot type, and the forgotten values refer to the missed annotations.

To solve those kinds of errors, Eric et al. (2019) adopted both manual corrections and automated corrections. After asking the human annotator to go over each dialog turn-by-turn, they also wrote scripts to canonicalize slot values to match the entities from the database. Besides, they also kept multiple slot values for over 250 turns in the case that multiple values are included in the dialog context. In addition to correcting dialog states, MultiWOZ 2.1 also corrects typos within dialog utterances for better research exploration of copy-based dialog models. As a result, over 40% of turns (around 30% of dialog state annotations) are corrected. Finally, MultiWOZ 2.1 also adds slot description for exploring few-shot learning and dialog act based on the pipeline from (Lee et al., 2019).

## 2.2 MultiWOZ 2.2

Building on version 2.1, MultiWOZ 2.2 (Zang et al., 2020) further proposes four remaining error types for dialog state annotations: early markups, annotations from database, more typos and implicit time processing. Apart from those semantic errors, MultiWOZ 2.2 also identifies the inconsistency of the dialog state annotations. For example, a slot value can be copied from dialog utterance, or derived from another slot, or generated based on the ontology. They also identify issues with the ontology, e.g., the format of values is not consistent and 21% of the values don’t match the database.

MultiWOZ 2.2 designed a schema to replace the original ontology and divided all slots into two types: categorical and non-categorical. For categorical slots, the possible slot values are limited and less than 50. Any value that is outside the scope of the database is labeled as “unknown”. On the other hand, values of non-categorical slots are directly copied from the dialog context and the slot span annotation is introduced to record the place and type of those non-categorical slots. Since typographical errors are inevitable in practice, MultiWOZ 2.2 leaves such errors in dialog utterances, hoping to train more robust models. In total, MultiWOZ 2.2 fixes around 17% of dialog state annotations, involving around 30% of dialogs.

In addition to the correction for the dialog states, MultiWOZ 2.2 also improve the annotations for the dialog acts. Though MultiWOZ 2.1 has added the dialog acts for the user side, there are still 5.82% of turns (8,333 turns including both user and system sides) lacking dialog act annotations. After employing crowdsourcing to complete the annotations, MultiWOZ 2.2 also renames the dialog acts by removing the prefix, so that the annotation of dialog acts can be used across all domains.

Our work builds on MultiWOZ 2.2, and further explores the annotation inconsistency (Section 3) and entity bias issues (Section 4) in the dataset.

## 3 Annotation Inconsistency

MultiWOZ is collected following the Wizard-of-Oz setup (Kelley, 1984), where each dialog is con-

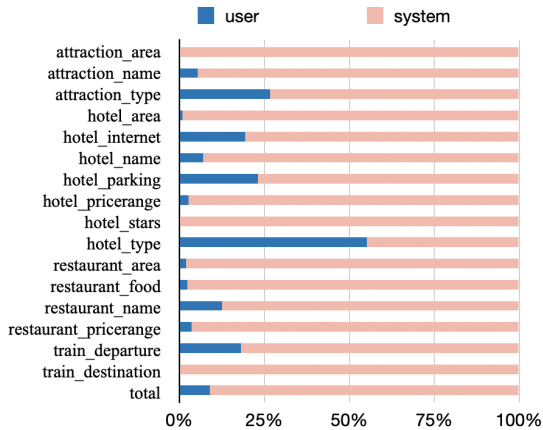


Figure 4: The proportion of whether the newly-added slots are extracted from user utterances (dark blue) or system responses (light red).

ducted by two crowd-workers. One crowd-worker plays the role of a human user and another one plays the role of the dialog system. The user would be assigned with a task goal, which describes the information of the task object and is required to conduct the conversation sticking to the task goal. The dialog system is required not only to respond to the user and complete the task, but also to take down the task-related information in the form of the slot values. Since the slots are annotated by the crowd-worker, and different dialogs employ different crowd-workers, the strategies they decide whether to take the information down are also different. Especially for the information provided by the dialog system, some crowd-workers decide to take it down while some do not.

In Figure 1, we show two dialogs with similar context. Both of the users ask for an **architectural attraction** in the **centre** area, and both of the dialog systems respond with a result, **all saints church**. Then, both users acknowledge the result and ask for more information on the result. However, in the left dialog, the attraction name, **all saints church**, is not annotated as one of the slot values, whereas the right dialog includes it in the dialog state. The source of this discrepancy may be that the annotator in the left dialog thinks the dialog system already knows this information and only information provided by the user should be annotated. On the other hand, the annotator in the right dialog might notice that the user has acknowledged the result, attraction name, and asked questions based on this result, hence it should be included in the dialog state. Having said that, the system cannot answer such follow-up questions without the annotation

Domain	Slot Type	Train	Valid	Test
attraction	area	491 (18.3%)	86 (21.2%)	76 (19.0%)
	name	1019 (38.0%)	151 (37.3%)	142 (35.6%)
	type	674 (25.1%)	102 (25.2%)	107 (26.8%)
	total	1773 (66.1%)	283 (69.9%)	256 (64.2%)
hotel	area	810 (24.0%)	120 (28.6%)	97 (24.6%)
	internet	657 (19.5%)	86 (20.5%)	75 (19.0%)
	name	1319 (39.1%)	166 (39.6%)	150 (38.0%)
	parking	638 (18.9%)	87 (20.8%)	71 (18.0%)
	pricerange	970 (28.8%)	114 (27.2%)	104 (26.3%)
	stars	665 (19.7%)	106 (25.3%)	91 (23.0%)
	type	1460 (43.3%)	195 (46.5%)	185 (46.8%)
total	2907 (86.2%)	360 (85.9%)	346 (87.6%)	
restaurant	area	799 (20.8%)	99 (22.0%)	82 (18.4%)
	food	689 (17.9%)	80 (17.8%)	88 (19.8%)
	name	1520 (39.6%)	189 (42.1%)	131 (29.4%)
	pricerange	792 (20.6%)	105 (23.4%)	82 (18.4%)
total	2635 (68.6%)	318 (70.8%)	257 (57.8%)	
taxi	departure	18 (1.2%)	1 (0.5%)	0 (0.0%)
	destination	14 (1.0%)	1 (0.5%)	1 (0.5%)
	total	31 (2.1%)	2 (0.9%)	1 (0.5%)
train	departure	70 (2.4%)	5 (1.0%)	8 (1.6%)
	destination	124 (4.2%)	19 (4.0%)	1 (0.2%)
	total	185 (6.2%)	23 (4.8%)	8 (1.6%)
total	5950 (74.2%)	768 (76.8%)	715 (71.5%)	

Table 1: Number and (percentage) of modified dialogs to correct dialog annotation inconsistencies.

of the attraction name in the dialog state. More examples of the inconsistency from other slot types or domains are listed in Figure 10 in Appendix.

After exploring all the dialogs from the test set, we manually examined and corrected each test dialog turn by adding the missing annotations. We first checked whether there are any missing annotations from the user utterance. For example, in the sixth turn of dialog “MUL0690.json”, user is asking for “*a moderate hotel with free wifi and parking*”. However, the token “*moderate*” from “*pricerange*” type is not included in the annotations. So, we added the slot “*hotel pricerange moderate*” to this turn. On the system side, for each dialog turn, we identified a possible slot value in the system response which is not included in the annotations, e.g., “*all saints church*” in the left dialog (“MUL1088.json”) in Figure 1, which is determined to be a possible slot value from “*attraction*” domain and slot type “*name*” based on the database and ontology file. Then, we examined those dialogs with slot annotations of “*attraction*” domain and slot type “*name*”. If we could find such dialogs with similar dialog context and containing annotations of the same domain and slot type, we complemented the annotations by adding the missing slot value.

In Figure 4, we illustrate the fraction of added annotations that come from the user utterance vs. system side. Each row corresponds to a slot type from a certain domain. As can be seen, the majority of the added utterances are from the system side,



Domain-Slot_Type	$H_1/H_0$	$H_\infty/H_0$
hotel-parking	0.217	0.060
hotel-internet	0.225	0.053
hotel-stars	0.592	0.249
restaurant-food	0.638	0.377
hotel-name	0.743	0.472
train-destination	0.753	0.269
hotel-stay	0.757	0.673
train-departure	0.776	0.288
attraction-area	0.792	0.355
train-leaveat	0.801	0.681
restaurant-area	0.824	0.384
restaurant-time	0.833	0.758
train-arriveby	0.850	0.732
attraction-type	0.852	0.514
attraction-name	0.855	0.636
restaurant-name	0.877	0.709
train-people	0.886	0.615
taxi-arriveby	0.890	0.736
taxi-departure	0.901	0.579
hospital-department	0.926	0.530
taxi-destination	0.936	0.685
taxi-leaveat	0.942	0.781
hotel-pricerange	0.944	0.662
hotel-area	0.954	0.658
hotel-type	0.969	0.729
restaurant-pricerange	0.971	0.746
train-day	0.999	0.947
hotel-day	0.999	0.954
restaurant-day	0.999	0.955
restaurant-people	0.999	0.969
hotel-people	0.999	0.973

Table 2: The unbalanced distribution among different slot types, measured using  $H_1/H_0$  (normalized Shannon entropy) and  $H_\infty/H_0$  (normalized min-entropy).

which confirms our original hypothesis: annotators often have no disagreement to take down information from the user utterance. However, they have different opinions about whether to annotate slots based on system responses.

For the training and validation sets, we write regular expressions that match the test set corrections and apply the scripts to automatically correct the annotations based on the database and ontology file, and modify the dialogs automatically. Table 1 list the corrected dialog numbers of each slot type and each domain, as well as the percentage of the corrected dialogs from all the dialogs in that domain. On average, about 20% of the dialogs involved slot modification for each slot type. The “name” and “type” slot types involve the most modification with around 40%. As we mainly focus on the missing slots extracted from system responses in these automated scripts, we ignore the slot types that can be solely modified by the user utterance, such as “book day” and “book people”. As can be seen, this process resulted in modification of totally more than 70% of the dialogs. To verify the correctness of the automated correction method, we randomly

sampled 100 modified dialogs and 100 unchanged dialogs, and check them manually. The verification result is shown in the table 3, based on which we compute the recall, precision and F1 score: 0.970, 0.961, 0.974.

	True	False
Positive	97	3
Negative	96	4

Table 3: Verification of the automated correction of the training/validation set.

## 4 Entity Bias

As discussed previously, another issue that we observe with MultiWOZ is the entity bias (e.g., “cambridge” appears in train destination in the majority of dialogs – Figure 2). Besides the “train-destination” slot type, we further explore the similar bias problem in all other 30 slot types. For each slot type, we quantify the frequency at which each possible slot value appears in the training data.

We quantify the entity bias using two metrics. For a vector  $r = (r_1, \dots, r_R)$  of frequencies of  $R$  entities, we define the normalized Shannon entropy as

$$H_1/H_0 := \sum_{i \in [R]} r_i \log_R \left( \frac{1}{r_i} \right). \quad (1)$$

Normalized Shannon entropy is bounded between 0 and 1, where  $H_1/H_0 = 0$  implies a deterministic distribution with all weight on a single entity and 1 implies a perfectly uniform distribution. Since Shannon entropy does not capture the tail of the distribution, we also report the normalized min-entropy, which is given by

$$H_\infty/H_0 := \max_{i \in [R]} \log_R \left( \frac{1}{r_i} \right). \quad (2)$$

Min-entropy captures the normalized likelihood of the most frequently appearing entity in the list of all possibilities. For example, as in Figure 2, frequency of the entity “cambridge” is about 50% which is much higher than 7% (which would have been its frequency had all 13 possible entities were uniformly distributed).

The entity bias for all 30 slots in the dataset is depicted in Table 2, ordered from the least uniform to the most uniform as measured by *normalized Shannon entropy*. We observe that some slot types, such as “hotel parking” and “hotel-internet” are significantly biased. There are only three possible slot

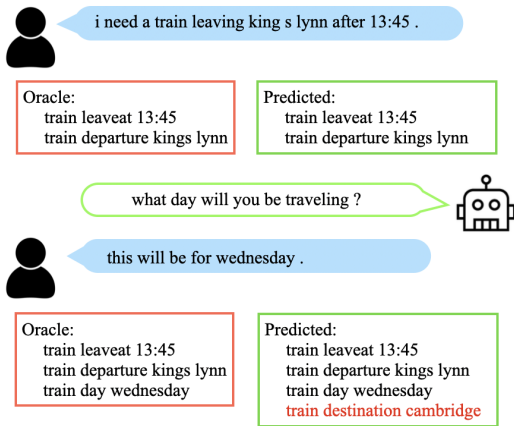


Figure 5: The first two turns from dialog “PUM1812” with generated dialog states from the SimpleTOD (Hosseini-Asl et al., 2020), which is trained on MultiWOZ 2.2. The dialog context does not mention “cambridge”, but the model generates this token.


values for the slot type “hotel-internet”: “yes”, “no” and “free”, where their count numbers are 10, 023, 326 and 9 correspondingly. We also find that many other slot types, such as “restaurant-food” suffer from severe entity bias besides “train-destination” as well.

On the other hand, the slot types involving day and people seem to be nicely balanced, such as “hotel-day” and “hotel-people”. This might be because those values are actually uniformly made up, while values of other slot types like “type”, “food” are real values and indeed follow certain real-world distributions.

These entity biases are potentially amplified by the learning models, which would lead to biased generation. In Figure 5, we show one such case from SimpleTOD (Hosseini-Asl et al., 2020) where in the current turn, the user is providing the information of “day” in the “train” domain. The dialog state tracking model successfully extracts the token “wednesday” and updates dialog states in the red rectangle. However, the model also adds the dialog state “train destination cambridge”, while “cambridge” has never been mentioned in the dialog history, which is potentially explained by the severe entity bias present in the “train-destination.”

Different from the annotation inconsistency problem, we do not make any modification to the training dataset based on our observation with respect to the entity bias. Strictly speaking, bias cannot be considered as a source of error in the dataset, and it needs to be tackled via better modeling efforts. Although the entity bias hurts the

Original dataset



After replacing entities




Figure 6: Example of dialog with new entities by replacing “bridge guest house” with “best western of long beach”

prediction accuracy of low-frequency slots and results in generating extra high-frequency slots, it also reflects certain real-world facts/biases as the dialogs are conducted by humans. This usually helps the learning task with limited training data, e.g., dialog domain adaptation (Qian and Yu, 2019; Lu et al., 2020).

While we keep the bias in the original dataset intact, we propose a new test set with all entities replaced with new ones unseen in the training data to facilitate the identification of whether models capitalize on such biases. For each slot type from each domain in the MultiWOZ, we find a similar slot type in the Schema-Guided dataset (Rastogi et al., 2019). For the slot values belonging to those slot type, we replace them with unseen values from the Schema-Guided dataset. Examples of dialog with replaced entities, along with predicted slots by our benchmark model is shown in Figure 6 and Figure 11 (Appendix).

## 5 Benchmarking State-of-the-Art Models

To verify our corrections of the dialog state annotations, we benchmark state-of-the-art dialog state tracking (DST) models on our modified dataset.

Traditionally, for DST task, the slot value is predicted by selecting from pre-defined candidates or extracting from dialog context. We adopt TRADE (Wu et al., 2019) as a representative of the mixture of these two methods. More recent works focus more on fine-tuning pre-trained model, which purely generates slot values based on dialog history. We choose SimpleTOD (Hosseini-Asl et al., 2020) and fine-tuned BART (Lewis et al.,

Models	Standard Results			Fuzzy Results		
	2.1	2.2	Ours	2.1	2.2	Ours
TRADE (Wu et al., 2019)	44.4±0.3	45.6±0.5	55.2±0.2	45.1±0.3	46.9±0.2	58.2±0.4
SimpleTOD (Hosseini-Asl et al., 2020)	54.7±0.5	53.6±1.0	62.1±0.2	55.2±0.5	54.4±1.2	64.7±0.2
DST-BART (Lewis et al., 2020)	57.9±0.5	56.0±0.7	67.4±0.5	58.7±0.2	57.5±0.2	72.3±0.4

Table 4: The performance of TRADE, SimpleTOD and DST-BART in terms of joint goal accuracy on MultiWOZ 2.1, 2.2 and our modified version. Fuzzy Results considers model-predicted slot value is correct if it is very similar with the ground truth even if they are not exactly the same.

2020) as benchmark models for DST as well.

**TRADE** (Wu et al., 2019) integrates GRU-based (Cho et al., 2014) encoder-decoder model and pointer-generator (See et al., 2017) to learn to copy slot values either from the dialog context or from the pre-defined value candidates.

**SimpleTOD** (Hosseini-Asl et al., 2020) builds a DST model by fine-tuning GPT2 (Radford et al., 2019), a large pre-trained language model. It combines all the condition information, including dialog history, previous dialog states and user utterance into a single sequence as input and let the language model learn to generate a sequence, containing dialog states and system response. Here we only feed in dialog states as ground truth output during training step, so that the trained model is specially designed for DST.

**DST-BART** builds a DST model by fine-tuning BART (Lewis et al., 2020) on the DST task in MultiWOZ. BART is a denoising autoencoder pre-trained with corrupted text, making it more robust to noisy data. It consists of a bidirectional encoder and a left-to-right autoregressive decoder.

**Joint Goal Accuracy.** We adopt joint goal accuracy and slot accuracy as our metrics of interest to evaluate the performance of the benchmark models. The joint goal accuracy measures the ratio of the dialog turns in the entire test set, where all the slots, in the form of triplets (domain, slot type, slot value) are predicted precisely correctly. Instead of checking every dialog turn, slot accuracy checks each slot individually for all slot types.

The evaluation results are reported in Table 4. As can be seen, the performance of all baselines increased about 7-10% where the largest jump happened in TRADE. We also observe that while DST-BART and SimpleTOD are within statistical error on MultiWOZ 2.2, DST-BART outperforms SimpleTOD by 5% points on our modified dataset. We suspect one potential reason is that BART is pre-trained on corrupted text, which improves its robustness and ability to handle noisy text. Another reason might be that BART consists of an encoder


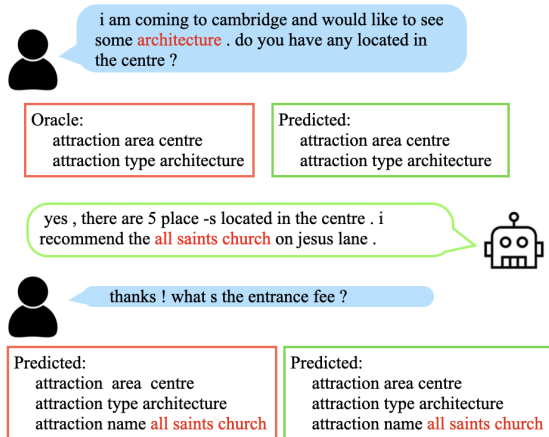


Figure 7: Fraction of error turns on the new corrected data.


and a decoder, containing 400M parameters, while GPT2 is a decoder only with 117M parameters.

The improvement from previous dataset versions comes mostly from the removal of the confusion because of the inconsistency. With the added slots, some of the predictions previously marked as erroneous are now recognized as correct. As shown in Fig. 8(a), the “attraction name all saints church” was not included in the ground truth in the old dataset version, so the prediction made by SimpleTOD was considered as wrong. With the modified dataset, SimpleTOD makes the same prediction, and since the annotation is more consistent, the models are more confident to learn the pattern and less likely to miss predicting slots. The changes in these two aspects leads to the drop of error turn numbers, shown in Fig 8(b), especially for the slot type “name” related, which involves the most slot modifications, corresponding to Table 1. The happens at the costs of slight increase of the error turn numbers for other slot types result from the increasing of the total turn number for those slot types, since we add slot annotations in the modified dataset. Overall, the percentages of turns with error all decrease, shown in the Fig. 7.

**Fuzzy Match.** There is another issue that multiple



(a) The generated result for dialog "MUL1088.json"



(b) Distribution of turn numbers with err over slot types

Figure 8: (a) The generated dialog state results (in green rectangles) using SimpleTOD. The model generates the “*attraction name all saints church*” in the second turn, however, the slot is included in the ground truth after the annotation correction; (b) The distribution of the numbers of turns where SimpleTOD makes a mistake, testing on MultiWOZ 2.2 test set (red bars) and our modified version (blue bars), showing a drop in errors associated with “*attraction-name*” and “*restaurant-name*”.

slot values can refer to the same item. For example, in the fifth turn of dialog “MUL0148.json”, the user and system are talking about booking at the “*huntingdon marriott hotel*”. The ground truth annotation for the hotel name is “*huntingdon marriott*”, while SimpleTOD predicts “*huntingdon marriott hotel*”, which is also the value in the dialog context. We believe these kinds of mismatches should be ignored (as they can be fixed via simple wrappers to find the closest match) and attention should be focused to other dominant problems to improve building greatly-performing DST models. As such, we adopt Levenshtein distance to compute the similarity between the ground truth and the predicted slot values. We consider the prediction to be correct if the similarity is above 90%. The result is listed in the Table 4. The performances after the fuzzy matching increases by 1-5%, consistent with the standard results.

	Joint Goal Acc.
MultiWOZ 2.1 test set	56.0±0.7
New test set with replaced entities	27.0±2.0

Table 5: Performance of DST-BART on MultiWOZ 2.1.

**New test set.** Similar with (Raghu et al., 2019), we evaluate SOTA models on a new test set, where we replace the slot entities with unseen values, resulting in a 29% performance drop in the terms of joint goal accuracy on DST-BART. In Figure 9, we show an example dialog, which was correctly predicted by DST-BART on the original test set. As can be seen, on the new test set, the model predicts entities

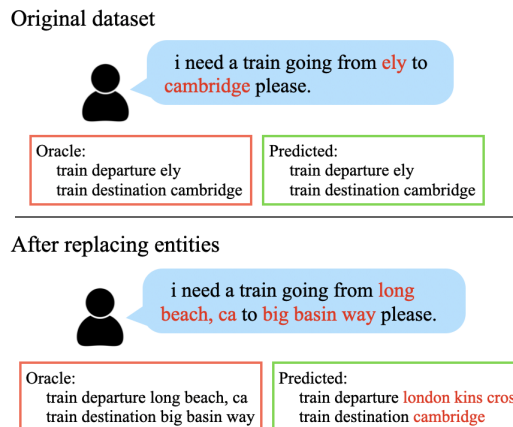


Figure 9: The outputs from DST-BART on the first turn of dialog “SNG0293.json”, before and after replacing the entities.

that never appear in the dialog context hinting at severe memorization of these named entities.

## 6 Concluding Remarks

MultiWOZ is a well-annotated task-oriented dialog dataset, and widely used to evaluate dialog-related tasks. Previous works like MultiWOZ 2.1 and MultiWOZ 2.2 have carefully identified and corrected several errors in the dataset, especially for the dialog state annotations. Building on MultiWOZ 2.2, we identified annotation inconsistency across different dialogs as a source of confusion for training dialog state tracking models. We proposed a correction and released a new version of the data with corrections. We also identified named entity bias as another source of issue, and released a new test set with all named entities replaced with unseen ones. Finally, we benchmarked a few state-of-the-art dia-



log state tracking models on the new versions of the data, showing 5-10% performance improvement on the new corrected data, and 29% performance drop when evaluation is done on the new test set with replaced entities. We hope the better understanding of MultiWOZ helps us gain more insights into dialog evaluation on this dataset.

While we corrected some errors in this work, we observe a few remaining problems in MultiWOZ. First, there are some cases where the annotation contradicts with the database. For example, in dialog “MUL2523.json”, the user is asking about “autumn house”, which is annotated as type of “guest house” in the database. However, the dialog state annotation labels the hotel type as “hotel”. This disagreement might not hurt the training of dialog state tracking models, but affects the now popular end-to-end dialog models, which are trained on MultiWOZ. There still some annotation errors. For example, in dialog “MUL0072.json”, the “monday” has never been mentioned, while the dialog states include the slot “hotel day monday”.

Finally, the source of the inconsistencies identified in this work is the Wizard of Oz data collection strategy, where different crowd-workers may annotate the dialogs differently. One way to mitigate such confusions might be to provide annotators with crystal clear annotation guidelines, or to have each dialog annotated by multiple annotators.

## References

- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Ultes Stefan, Ramadan Osman, and Milica Gašić. 2018. Multiwoz - a large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Bill Byrne, Karthik Krishnamoorthi, Chinnadhurai Sankar, Arvind Neelakantan, Ben Goodrich, Daniel Duckworth, Semih Yavuz, Amit Dubey, Kyu-Young Kim, and Andy Cedilnik. 2019. *Taskmaster-1: Toward a realistic and diverse dialog dataset*. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4516–4525, Hong Kong, China. Association for Computational Linguistics.
- Kyunghyun Cho, B. V. Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. On the properties of neural machine translation: Encoder-decoder approaches. In *SSST@EMNLP*.
- Mihail Eric, Rahul Goel, Shachi Paul, Abhishek Sethi, Sanchit Agarwal, Shuyag Gao, and Dilek Hakkani-Tur. 2019. Multiwoz 2.1: Multi-domain dialogue state corrections and state tracking baselines. *arXiv preprint arXiv:1907.01669*.
- Ting Han, Ximing Liu, Ryuichi Takanobu, Yixin Lian, Chongxuan Huang, Wei Peng, and Minlie Huang. 2020. Multiwoz 2.3: A multi-domain task-oriented dataset enhanced with annotation corrections and co-reference annotation. *arXiv preprint arXiv:2010.05594*.
- M. Heck, Carel van Niekerk, Nurul Lubis, Christian Geischauser, Hsien-Chin Lin, M. Moresi, and Milica Gavsic. 2020. Trippy: A triple copy strategy for value independent neural dialog state tracking. In *SIGdial*.
- Ehsan Hosseini-Asl, B. McCann, Chien-Sheng Wu, Semih Yavuz, and R. Socher. 2020. A simple language model for task-oriented dialogue. *ArXiv*, abs/2005.00796.
- J. F. Kelley. 1984. An iterative design methodology for user-friendly natural language office information applications. *ACM Trans. Inf. Syst.*, 2:26–41.
- Sungjin Lee, Qi Zhu, Ryuichi Takanobu, Xiang Li, Y. Zhang, Zheng Zhang, Jin chao Li, Baolin Peng, Xiujun Li, Minlie Huang, and Jianfeng Gao. 2019. Convlab: Multi-domain end-to-end dialog system platform. In *ACL*.
- M. Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, A. Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2020. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *ArXiv*, abs/1910.13461.
- J. Lu, Pinghua Gong, Jieping Ye, and C. Zhang. 2020. Learning from very few samples: A survey. *ArXiv*, abs/2009.02653.
- Baolin Peng, C. Li, Jin chao Li, Shahin Shayandeh, L. Liden, and Jianfeng Gao. 2020. Soloist: Few-shot task-oriented dialog with a single pre-trained auto-regressive model. *ArXiv*, abs/2005.05298.
- Kun Qian and Z. Yu. 2019. Domain adaptive dialog generation via meta learning. In *ACL*.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Dinesh Raghu, Nikhil Gupta, and Mausam. 2019. *Disentangling Language and Knowledge in Task-Oriented Dialogs*. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1239–1255, Minneapolis, Minnesota. Association for Computational Linguistics.

- Abhinav Rastogi, Xiaoxue Zang, Srinivas Sunkara, Raghav Gupta, and Pranav Khaitan. 2019. Towards scalable multi-domain conversational agents: The schema-guided dialogue dataset. *arXiv preprint arXiv:1909.05855*.
- L. Rojas-Barahona, M. Gaić, N. Mrksic, P. Su, Stefan Ultes, Tsung-Hsien Wen, S. Young, and David Vandyke. 2017. A network-based end-to-end trainable task-oriented dialogue system. In *EACL*.
- A. See, Peter J. Liu, and Christopher D. Manning. 2017. Get to the point: Summarization with pointer-generator networks. *ArXiv*, abs/1704.04368.
- Jianhong Wang, Yeliang Zhang, Tae-Kyun Kim, and Yunjie Gu. 2020a. Modelling hierarchical structure between dialogue policy and natural language generator with option framework for task-oriented dialogue system. *ArXiv*, abs/2006.06814.
- Kai Wang, Jun-Feng Tian, Rui Wang, Xiaojun Quan, and J. Yu. 2020b. Multi-domain dialogue acts and response co-generation. *ACL 2020*.
- Chien-Sheng Wu, Andrea Madotto, Ehsan Hosseini-Asl, Caiming Xiong, Richard Socher, and Pascale Fung. 2019. Transferable multi-domain state generator for task-oriented dialogue systems. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics.
- Qing yang Wu, Yichi Zhang, Yu Li, and Z. Yu. 2019. Alternating recurrent dialog model with large-scale pre-trained language models. *ArXiv*, abs/1910.03756.
- Xiaoxue Zang, Abhinav Rastogi, Srinivas Sunkara, Raghav Gupta, Jianguo Zhang, and Jindong Chen. 2020. Multiwoz 2.2: A dialogue dataset with additional annotation corrections and state tracking baselines. In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI, ACL 2020*, pages 109–117.
- Jian'guo Zhang, Kazuma Hashimoto, Chien-Sheng Wu, Yao Wan, Philip S. Yu, R. Socher, and Caiming Xiong. 2019. Find or classify? dual strategy for slot-value predictions on multi-domain dialog state tracking. *ArXiv*, abs/1910.03544.
- Yichi Zhang, Zhijian Ou, Huixin Wang, and Junlan Feng. 2020a. A probabilistic end-to-end task-oriented dialog model with latent belief states towards semi-supervised learning. *ArXiv*, abs/2009.08115.
- Yichi Zhang, Zhijian Ou, and Z. Yu. 2020b. Task-oriented dialog systems that consider multiple appropriate responses under the same context. *ArXiv*, abs/1911.10484.

<p><b>Restaurant-Name</b></p> <p>Dialog: MUL0011.json, Turn: 3 Context: &lt;system&gt; the <b>clowns cafe</b> is in the centre area. would you like me to book you a reservation? &lt;user&gt; yes, please. i'd like a table for 6 at 18:45 on friday .</p> <p><b>With slot:</b> restaurant name <b>clowns cafe</b></p>	<p>Dialog: MUL0011.json, Turn: 2 Context: &lt;system&gt; <b>clowns cafe</b> is a restaurant in the city centre with that price range. would you like to make a reservation? &lt;user&gt; sure! how close is the cafe from my current location?</p> <p><b>Without slot:</b> restaurant name <b>clowns cafe</b></p>
<p><b>Hotel-Name</b></p> <p>Dialog: PMUL0204.json, Turn: 8 Context: &lt;system&gt; yes. the <b>lensfield hotel</b> is located in the south part of town. would you like more information or to book it? &lt;user&gt; does it have free parking?</p> <p><b>With slot:</b> hotel name <b>lensfield hotel</b></p>	<p>Dialog: MUL0021.json, Turn: 2 Context: &lt;system&gt; the <b>lensfield hotel</b> is the only hotel that matches your criteria. would you like to book a stay? &lt;user&gt; i'd just need their star rating and phone number, thank you.</p> <p><b>Without slot:</b> hotel name <b>lensfield hotel</b></p>
<p><b>Hotel-Type</b></p> <p>Dialog: MUL1088.json, Turn: 2 Context: &lt;system&gt; there is no entrance fee, it is free. is there anything else i can help you with? &lt;user&gt; i'd like to find a <b>guesthouse</b> to stay at. i want it to be nice, so i'd like an expensive one.</p> <p><b>With slot:</b> hotel type <b>guesthouse</b></p>	<p>Dialog: MUL2009.json, Turn: 3 Context: &lt;system&gt; do you need help booking a ticket on that train or with anything else? &lt;user&gt; i want to make a booking for 7 people, and i need the reference number. also, i need to book an expensive <b>guesthouse</b>.</p> <p><b>With slot:</b> hotel type <b>guesthouse</b></p>

Figure 10: More Examples of dialog with or without slot annotations for different slot types.

domain	slot type	train	valid	test
attraction	area	2081 (13.18%)	367 (15.17%)	350 (14.33%)
	name	3738 (23.67%)	558 (23.06%)	505 (20.67%)
	type	2915 (18.46%)	426 (17.60%)	468 (19.16%)
	total	7288 (46.15%)	1163 (48.06%)	1078 (44.13%)
hotel	area	3486 (15.72%)	530 (18.96%)	454 (17.37%)
	internet	3175 (14.32%)	400 (14.31%)	340 (13.01%)
	name	3998 (18.03%)	616 (22.03%)	479 (18.32%)
	parking	3122 (14.08%)	431 (15.41%)	326 (12.47%)
	pricerange	4342 (19.59%)	485 (17.35%)	443 (16.95%)
	stars	2989 (13.48%)	472 (16.88%)	406 (15.53%)
	type	7061 (31.85%)	985 (35.23%)	919 (35.16%)
total	15048 (67.88%)	1929 (68.99%)	1822 (69.70%)	
restaurant	area	2906 (12.57%)	435 (14.99%)	352 (12.21%)
	food	2648 (11.45%)	361 (12.44%)	349 (12.11%)
	name	4705 (20.35%)	702 (24.20%)	458 (15.89%)
	pricerange	3036 (13.13%)	441 (15.20%)	367 (12.73%)
	total	9932 (42.95%)	1363 (46.98%)	1051 (36.46%)
taxi	departure	32 (0.70%)	2 (0.29%)	0 (0.00%)
	destination	23 (0.50%)	1 (0.15%)	1 (0.16%)
	total	53 (1.16%)	3 (0.44%)	1 (0.16%)
train	departure	235 (1.29%)	14 (0.48%)	35 (1.19%)
	destination	441 (2.43%)	61 (2.10%)	5 (0.17%)
	total	658 (3.62%)	75 (2.59%)	35 (1.19%)
total	27656 (50.50%)	3783 (51.92%)	3467 (47.68%)	

Table 6: Number of dialogs involving modification for inconsistency distribution.

<p>Context (‘MUL1901.json’, turn 12)</p>	<p>&lt;user&gt; i need a place to stay &lt;system&gt; is there a certain area or price range you had in mind? &lt;user&gt; i really want to stay with something moderately priced. but, i want it to be 4 star rated if possible. &lt;system&gt; there are lots of moderately priced four star hotels. are you interested in a particular part of town? &lt;user&gt; maybe a guesthouse with free parking, i have 7 guests and that will be for 2 nights this monday &lt;system&gt; we have several such choices in the north, south and east. if you have no preference for location shall i just book one for you? &lt;user&gt; that would be great. &lt;system&gt; <b>hotel kings cross</b> is a moderate price, 4 star guesthouse in the east area with wifi and parking. will that work? &lt;user&gt; yeah, that sounds great. do they have room for all 7 of us? &lt;system&gt; i can ... from monday. &lt;system&gt; done! your reference number is 0c3z1qko . can i help you with anything else? &lt;user&gt; i am also looking to visit cinelux almaden cinema, do you know the address and cost to visit?</p>
<p>Ground Truth</p>	<p>hotel bookstay 2, hotel parking yes, hotel bookpeople 7, hotel type guesthouse, hotel bookday monday, hotel stars 4, hotel pricerange moderate, hotel name <b>hotel kings cross</b>, attraction name saint johns college, attraction name cinelux almaden cinema</p>
<p>Model(BART) Result</p>	<p>hotel bookstay 2, hotel parking yes, hotel pricerange moderate, attraction name cinelux almaden cinema, hotel name <b>a and b guest house</b>, hotel bookday monday, hotel type guesthouse, hotel stars 4, hotel bookpeople 7</p>

Figure 11: Example of dialogs with new entities